"BBN Job #41540"
Len Kleinrock
Tape Number 5
<u>-- **CARIBINER GROUP**</u>

(OFF MIKE)

QUESTION

LEN KLEINROCK

First time I heard about the ARPANET was when Larry Roberts(?) approached me and said, Len, I've got this great idea and this great need. We've got to connect computers together through a network for a variety of important reasons. His main reason was that he wouldn't have to buy a new computer for every researcher that he supported in ARPA(?). So the original idea of the ARPANET was to do this sharing of resources. And it wasn't a military application that he told us.

I don't know what he told the Pentagon, but his ... his ... his idea was, he wouldn't have

to buy everybody a computer, and those that he did buy computers were doing special things with them. Example, at Utah they were doing special graphics, UCLA we were doing simulation packages, uh, to SRI(?) they were doing information and storage and retrieval. And these became very unique resources, there was a super computer being developed at Illinois.

So he said, look. These unique resources need to be shared over the network, and I'm tired of buying every new re, investigator a new computer, so let's connect them together. That was the first knowledge that there was gonna be a thing called the ARPANET created.

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

Well, let's begin at the beginning. From my viewpoint, Phil(?) began with a comic book. At age six, I was reading the "Superman" comic book, and in the middle, it described how to build a crystal radio. What I needed was my father's old razor blade, a piece of pencil lead, an empty toilet paper roll, some wire. All of those I could get easily. I also needed an earphone, which I quickly appropriated from the public telephone company.

The last thing was something called a variable capacitor. Well, I got my mother to take me on the subway down to Canal Street, where all the radio shops are. I walk into the store, and I boldly ask the guy, I need a variable capacitor. And he said, what size? (Laughs) And that blew my cover. I had no idea. So I admitted what I was trying to do, he outfitted me with just the right part, went home, put it together, and magic happened. I got free music over the radio. No

batteries, no power, no nothing, it was free, and in some sense, an engineer was born that day.

Spent the rest of my youth cannibalizing discarded radios, putting together radios, designing them, building them. Went to Bronx Science, took the usual courses and studied radio engineering, and then I went off to City College. Couldn't afford to go, by the way, to a day session. Went to evening session, got a bachelor's degree in electrical engineering, was elected student body president of the evening session and got a scholarship to, full scholarship to MIT.

And that's when this networking side of my life begins, with the strong electrical engineering background.

QUESTION

LEN KLEINROCK

Not ...

QUESTION

LEN KLEINROCK

At MTI, some interesting things were going on. The first thing you do as a Ph.D. student ...

QUESTION

LEN KLEINROCK

Oh.

QUESTION

LEN KLEINROCK

In 1959, I began a Ph.D. program in the electrical engineering department. Now all of my class mates were busy refining a highly specialized area in communication theory, known as coding theory. And I decided that's not for me, it were just too little, too late. I wanted to break new ground. And so I began to work on this thing called a message switching communications network, and I began to study what later would become the notion of packet switching.

I studied the principles, the analysis, the design, mathematical treatment, algorithms, et

cetera. And around that time, I met Larry Roberts, he was a classmate of mine. He was doing something totally different, he was working on three-dimensional, uh, picture processing. And I also met Ivan Sutherland(?) at that time, a name that may or may not be ... be known to this group. Um, a key guy.

Anyway, Larry and I were strong classmates, and I was doing my work on networking and he was doing his work on picture processing. And, uh, I graduated, my work was finished in December '62, it became a book, by McGraw-Hill, Lincoln Laboratories, MIT, in 1964, and laid down the foundations for what later became the packet switching technology. Unfortunately, there was absolutely no one interested in this technology in 1962, '63.

In '63 I joined the faculty at UCLA, and continued to pursue the research, doing performance evaluation and design work, and

then in 1967, Larry approached me and said, Len, I've got this need. Every time, he was now working for ARPA, every time a new investigator wants to be funded, he asks me to buy them a computer. And I'm getting sick of buying, (Laughs), all these computers cause they're all the same machine. And every time they get one of their machines, they make something different out of it.

One guy puts graphics on it, you put a lot of simulation on it. They were doing artificial intelligence at MIT and Stanford. And now these unique resources are out there and I can't duplicate all that capability at every location. And I know you worked on networking. So let's put together a network which allows these various computers to be shared through the network, so you could use my simulation package and I can use your artificial intelligence package.

So in 1967, Larry gathered around him a small handful of people, including myself, to lay out the specification for what was later to become the ARPANET. I remember very clearly there was, um, Herb Baskin(?) was there, he had done some work in time sharing. And he banged his fist on the table and said, if this network can't give me a one second response time, it's no good for time sharing. And so we put down a spec, (Claps Hands), one second response time.

I was there, and I said, look, this is an experimental network. If you can't measure what's going on, you've got nothing. So we installed sophisticated requirements for measurement hooks. And so the, slowly, the spec began to evolve. Spec was written in mid-1968, a request for quotation(?) was sent out to industry, a number of industry organizations responded, BBN being one of them, and many of them bid the same machine.

What they were gonna do is, take a mini-computer, change some of the hardware and the software, make it into a machine that could do packet switching in a network. BBN won the contract in nineteen-seven ... sixty-nine, January of '69. And they were supposed to deliver to UCLA the first IMP, Interface Message Processor, the day after Labor Day in 1969. They chose UCLA as the first site because of my strong development of the early technology and my work all those years. That was the game plan.

And our job at UCLA was to be able to receive this switch and connect it to our host computer. Now that made us the first node(?) on the network. You might laugh, how can you be the first node on a network? And the answer is, you got to make those two machines talk to each other. When(?) we heard that BBN was two weeks late in delivering this machine, much

to our delight, cause we needed that time to get the spec ready(?). In fact, BBN wouldn't tell us how to talk to their machine, and yet, when they delivered the machine they expected us to be able to communicate.

So we had to squeeze out of them, plead out of them, the spec, invent part of the spec ourselves, write the software, build the hardware, et cetera. Machine's gonna be two weeks late, what could be better? Except they then ... then informed us they put the damn thing on an airplane, and it was gonna arrive on time. We worked day and night, we were ready when it arrived, it came in on a Tuesday, the day after Labor Day. You got a picture of the scene, okay? My guys were there. I had put together this team ...

QUESTION

LEN KLEINROCK

There is confusion there.

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

The early ideas of packet switching came from a number of sources independently. I in my Ph.D. dissertation, had laid down the mathematical principles, the design, the protocol evaluation, et cetera. The same time, Paul Barron(?) at Rand Corporation, had written some, was writing some reports in the early Sixties about a thing which amounted to packet switching as well. We both had the same idea of creating a mesh network, with dynamic routing through a network, with addressed packets moving through, highly reliable kind of system.

I did the analysis, at the same time, Paul was inventing similar ideas. Meanwhile, at the National Physical Laboratory, in, uh, in, uh, it

was New, was ... wasn't Newcastle, it was, remind me of the city somebody. Um ...

(OFF MIKE)

QUESTION

LEN KLEINROCK

The origins of packet switching came from more than one source. In my early work from '59 to '62, I had laid out the mathematical foundation for the way packet switching networks would behave and should be designed. There were mathematical models, simulation models, protocol models. In the early Sixties, Paul Barron produced some reports which described similar ideas from a, from a, uh, behavioral point of view.

He talked about the idea of random routing, as I had. He talked about packets, he talked about dynamic routing, he talked about mesh networks as ... as had I. Meanwhile, Donald Davies(?) at National Physical Laboratory

in England, was thinking about the same things, and he reached a stage of actually building a packet switch, even before the ARPANET was developed. But he never got beyond the one switch. Basically, they ran out of funding in England.

The same time, Larry Roberts, who had been a classmate of mine, in the mid-Sixties, joined the ARPA office, and he earlier had tried to make a connection from a machine at Lincoln Laboratory, to a machine at System Development Corporation, that's from Cambridge to Santa Monica. And he found that so painful, that he marked it in his head, some day he's gonna try to fix that. So here we have in the early Sixties, these pioneers developing technology of packet switching on paper.

Larry's at the ARPA office, and what is he doing, he's supporting research in time sharing. So along comes a new principal

investigator to him, Larry says listen, I'll support you. The investigator says terrific, buy me a computer, if you want me to do research on computers. Now, this was getting a little bit boring for Larry and expensive. He couldn't afford to give everybody a computer. So that, with his knowledge that there was the ability to create a network where these machines could be accessed remotely, led him to the concept of an ARPANET.

(OFF MIKE)

QUESTION

LEN KLEINROCK

Let's talk about where the concept of packet switching came from. In fact, there were a number of independent studies that led to this idea of packet switching. In '59, I began studying the principles, the mathematical analysis, the simulation and design technology for what became packet switching, it was in the

form of message switching at the time. And in '62, I presented my Ph.D. dissertation, which became a book published by, uh, MIT-Lincoln Laboratories.

The same time in the mid, the early Sixties, Paul Barron at Rand Corporation was conceiving similar ideas. These ideas between us were the idea of a mesh network with packets, blocks of data moving through a network independently, with dynamic routing, going around blocked and congested regions, et cetera. Highly robust, highly dynamic, very effective technology.

Meanwhile, in England, Donald Davies was conceiving of the notion of packet switching as well. And he went so far as to build a packet switch, a single note(?) packet switch before the ARPANET actually came on line. Unfortunately, he ran out of funding, and they couldn't pursue that project. But here we have these ideas

coming out of a number of different independent sources.

Meanwhile, I had met Larry at ... at MIT, he and I were very close classmates, in fact, we shared an office at Lincoln Laboratory together. Larry in the mid-Sixties, joined the ARPA office when Robert Taylor, who had been funding computer research in the Sixties, brought Larry in to put together the notion of a network. Bob Taylor basically sold the idea to the government, and Larry was there to make it happen.

Now, why did they want a network like this? Well, in the mid-Sixties the ARPA computer office was supporting the hottest topic around then, which was time sharing. And they went to a number of principal investigators and said listen, you want to do research on time sharing? The guy said yes. If you're gonna make me do research, buy me a computer. And

they kept buying computer after computer, and at each site, the computer changed form.

For example, at UCLA, we, and ... at UCLA, we were able to add very sophisticated simulation capabilities on our machine. At Utah they were adding graphics capability. At Carnegie-Mellon, MIT, Stanford, they were adding artificial intelligence. At Illinois, they were building a super computer. So all these unique resources were appearing throughout the network, and every time a new research ... researcher came on, there's no way that the ARPA office could bring all this capability to each site.

So Larry, having experimented with networking earlier, he had tried to connect a machine from Lincoln Laboratory in Cambridge, Mass., to a machine in Santa Monica at the System Development Corporation. And that was such an unpleasant, difficult procedure for him,

he noted in his head that some day, he'd like to see that done in a more sophisticated way.

So here we have, Larry who knows my work in networking, he recognized there's a need for networking sometime. He suddenly finds a need at the ARPA office, they're spending too much money on providing everybody with the same kind of computer. So let's create a network and make all these resources accessible to people through the network. Fine. (Claps Hands) Larry approached me in 1967 and he said, Len, join me with a few other people and let's write down the specifications to what a packet switching network should look like.

So we went there, and we laid out the spec. I remember very well Herb Baskin was there, and he had done some work on time sharing. And he pounded his fist on the table and said, look, if that network can't deliver a response time of one second, I can't use it for

time sharing, it's too slow. So we wrote down a spec, one second response time.

And I was there and I said, look, if you can't put measurement hooks(?) into this experimental network you'll never know what's going on, so we specified a very sophisticated measurement package, which indeed found it's way into the software of the network. So the spec evolved. In 1968 it was formalized and sent out for request for quotation, a number of organizations responded. And many of them responded with the same machine.

See, the idea of this proposal was to take an existing mini-computer, add some hardware, add some software, and make it behave like a packet switch. Many of these organizations selected the Honeywell DDP-516(?) mini-computer, and BBN was one of those. BBN won the contract. In January of '69 it was awarded to them. And their task was to deliver

the first IMP at UCLA the day before Labor Day in 1969. Turned out to be a Saturday before the Monday of Labor Day.

Meanwhile at UCLA it was our job to build the hardware and the software, to allow our host computer and their IMP to talk to each other. Now that sounds like a straightforward process. That makes us the first node on the ARPANET, which sounds like an oxymoron, how can you be the first node on the network? And the answer is, to get these two machines to talk to each other, was a really sophisticated task. Especially so because BBN would not release the IMP to host specification to us.

And we had to squeeze blood to get it out of them, and in some cases we couldn't and we invented our own specification. So finally we got this back, we were busy trying to implement. Recognizing that the switch is going to appear on a Saturday just before Labor Day. In mid-

August, we learned that BBN had slipped the schedule and they were behind two weeks and we were totally thrilled with that. We needed the time to continue our implementation.

We learned just before Labor Day that the sons of guns, they put the damn thing on an airplane and were shipping it out to us. That kept us up, us up a few nights, finishing up the specification, getting implementation. The machine arrived, and we were ready. Now the scene when that machine arrived is something to behold. Everybody who could be there was there.

First of all, me and my team were there. We had the people from the UCLA administration there. I mean, after all, you're gonna use our university? BBN was there. ARPA was there. Of course Honeywell was there. The people who built my host machine, the system, um, Scientific Data Systems, was there. We had

ATT Long Lines there, we were gonna use their lines. We had GTE, we were gonna use their local access. Everybody that could be there was there, and they're all ready to point the accusing finger at the other guy, if it didn't work.

And fortunately, bits were flowing that first day. the interface was a success. And by the following Tuesday, we had messages moving back and forth. Fortunately, it was a big success and that was the birth of the ARPANET. It happened right then, around Labor Day, 1969. It was a great success.

QUESTION

LEN KLEINROCK

Well I feel very stiff, frankly, but ...

(OFF MIKE)

(CUT)

(OFF MIKE)

QUESTION

LEN KLEINROCK

Okay. You might wonder who or where, what was really happening in terms of the impact, the long range impact. The answer's absolutely not. For example, BBN structured the machine so it could not support more than sixty-four computers ever. And later (Inaudible) had to expand that. Well, you know, now we have three million computers attached to the network, growing at over eighty percent a year. Nobody had any concept of the impact of this development at all.

You know, there's another story about this ... this IMP that arrived. This B ... DDP-516 was a machine that Honeywell featured in 1968, and that's why many people bid it. Well I remember seeing that machine and I believe it was, either the string(?) or the four joint computer conference in one of the major cities, probably Las Vegas. And they had this machine, it was a hardened(?) military version, hoisted up,

hanging by it's hooks, swinging and running. And there was some son of a gun with a sledge hammer whacking the damn thing, proving that it could continue to operate while they beat the hell out of it.

I believe honestly that that's the physical machine that, (Laughs), BBN delivered to us, and sure enough, it continued to work. And as you know, it's still at UCLA today. Decommissioned, but proudly standing there as an epic(?) of the past.

QUESTION

LEN KLEINROCK

Can I say something mean about them then?

(OFF MIKE)

LEN KLEINROCK

I mean, I'm ready to strangle my secretary too. There's a lot, there were, there were lots of screw up's there.

(OFF MIKE)

QUESTION

LEN KLEINROCK

The technology behind packet switching rests on some very sophisticated mathematical structures. And I hate to say it, but it's ... it's related to something called queueing(?) theory. Now, the word queueing can be spelled one of two ways, and I prefer to spell it Q-U-E-U-E-I-N-G, because you get five vowels in a row, it's the only word in English I know with five vowels in a row.

Queueing theory is a, is a theory which talks about how long we spend waiting in line. Now what the devil can that have to do with networking? Well the answer is this. When these packets, these blocks of data move through the network, they reach a switch and maybe somebody is being transmitted out of that switch when they want to go, so they've got to wait in line. So queues build up, delays get included, and if you listen to the next words I'm

gonna use, you'll see why it's important.

We talk about response time, throughput, blocking, capacity, delays, those are exactly the measures you need to evaluate how a network's performing. When I send a message through a network, I want to know how long is it ... it gonna take to get there? How many bits per second can I pump in? What is the capacity of this network in the first place? So these issues have to be dealt with with a rather sophisticated theory, and queueing theory is exactly the right mathematical structure.

For my own dissertation, I used that theory and basically enhanced it, so it would apply to packet switching. And there's some very sophisticated things that have to be done, but that's why this weird notion of queueing theory enters in the first place. You've got to be able to predict the way a network's gonna perform before you build it. You've got to be

able to decide how to design a network. In order to do so, you have to evaluate way, the way different designs will perform.

And so it's a rather sophisticated mathematical structure. You could do it analytically, you could do it by simulation, you can build the damn thing and then measure it and then it's too late. We chose to do the analytic approach, and the simulation approach.

QUESTION

LEN KLEINROCK

Okay.

QUESTION

LEN KLEINROCK

In 1966, Howie Frank, who was then an assistant professor up at Berkeley, and I was a, an assistant professor down to UCLA, called me and asked me to participate in a short course he was running up there. He had heard of me because I had published a book called

"Communication Networks", that was the name of my thesis. By the way, that made him pretty mad, I understand, cause he had hoped he could use that title for a book he later wrote, and he's constantly complaining that I stole the ... the only right name for a book of this sort.

Now Howie at that time was involved in networking things as well. He was involved with topological design procedures, where to put lines in a network. What kind of a topological structure do you want? I was concerned with the related issues. I would take a problem, for example, if you give me a network with a certain topology, how will it perform? And if you change the topology, how will it perform? Howie was more focused on showing what the topology should look like.

Well. So I was scheduled to give this short course, (Laughs), at Berkeley, but I was plan, I had planned a vacation in, uh, Sequoia,

with my family, and there I was, with my family in Sequoia and the plan was, I'd drive out of Sequoia early in the morning, go to Berkeley, give me lecture and come home. Well I had to wake up at five ... five a.m. in the morning, I had about three days of beard, a plaid shirt, straw hat and all my clothes and notes thrown in a little suit-carry over my shoulder. Get in my old station wagon and I'm driving out of Sequoia National Park, when I notice smoke behind me.

My brakes had totally failed. I was now doing forty-five miles an hour, approaching a ninety-degree curve. This was to be the end. There was no, (Laughs), I'd get out of this. So I hit the brakes, nothing happened. Tried the emergent(sic), nothing happened. (Claps Hands) Put it in reverse, nothing happened. (Claps Hands) Tried to turn the engine off, nothing happened. I opened the door, was about to jump out the left side, but there was a steep cliff so

that would have been the end, so I closed the door, put the belt back on.

I would have rolled it on the right side, but it was a, again, not a stiff wall but it would turn the car over. There I am. What would you do in a case like this? I did the only thing I knew left, I yelled help. (Laughs) And believe it or not, help came. Suddenly my rear left tire blew and it caused the car to slowly grind to a stop. I'll tell you later what happened, why that occurred. But meanwhile, I got to get to Howie. Now I need money to repair the car, and he was paying me some good money to give this lecture.

So I hop out, a park ranger comes by, he says, yep, your car stopped. I said, good. Take me out of here. So he took me to a garage, I told the garage guys, fix the car, and I started hitchhiking out of Sequoia. Ranger comes back and he says, you're not allowed to hitchhike, he throws me out of the park. So I'm at the exit of

the park, hitchhiking, looking like a bum and nobody's gonna pick me up.

Finally I get a ride, racing down to the airport in ... in ... in, uh, Fresno. Mad rush, run through the airport with my beard, my suit carrier, hop on the plane, just made within a minute. Get on the plane, we taxi out to the runway, plane stops. And the captain says, folks, we have a problem, our brakes aren't working. Takes us back to the airport, hop on any plane that's in Fresno, just to get out of Fresno, and this plane was going to Los Angeles. Hop on the plane, to make a long story short, I got there three hours late. Howie was amazed I got there at all.

(TAPE SPEED VARIES FROM THIS POINT ON)

LEN KLEINROCK

Paul Barron was in the audience waiting to hear me talk. He took the first three hours. He was

gonna talk the next day. I did the afternoon and I came back a few days later to finish up my half. Got back to my tent in Sequoia that night, and my wife casually says to me, did you have a good day? (Laughs)

QUESTION

LEN KLEINROCK

Okay.

QUESTION

LEN KLEINROCK

Okay. What happened between Howie Frank and myself was we met in this 1966 lecture. Shortly after that, as you know, Larry approached me to put together the ARPANET. And as we began to grow the ARPANET, you know, to connect four nodes together is easy, there's no topological design. When you start adding ten, twenty, thirty nodes, it becomes a problem. And I was aware that Howie Frank had the expertise to topological design.

I introduced him to Larry and Larry hired him on as a sub-contractor, to now take over the topological design of the network. Which then BBN would deploy. Howie's approach, as I said, was topological. Mine was analytic and delay oriented, and we built our different approaches, which tended to lead to the same kinds of design. So Howie, Larry and I worked very well together, it was a very good mix. Bob Kahn(?) was also, of course, making his presence known then, through his efforts at BBN, and as you know, later on, Howie Frank, Bob Kahn and myself wrote this classic paper, which described our experiences in putting the network together.

(OFF MIKE)

(END OF TAPE #5)